

Wiretapping a hidden network ^{*}

Haris Aziz², Oded Lachish¹, Mike Paterson¹, and Rahul Savani³

¹ Department of Computer Science, University of Warwick, CV4 7AL Coventry, UK
`{oded,msp}@dcs.warwick.ac.uk`

² Institut für Informatik, Universität München, 80538 München, Germany
`aziz@tcs.ifi.lmu.de`

³ Department of Computer Science, University of Liverpool, L69 3BX Liverpool, UK
`rahul.savani@liverpool.ac.uk`

Abstract. We consider the problem of maximizing the probability of hitting a strategically chosen hidden *virtual network* by placing a wiretap on a single link of a communication network. This can be seen as a two-player win-lose (zero-sum) game that we call the *wiretap game*. The *value* of this game is the greatest probability that the wiretapper can secure for hitting the virtual network. The value is shown to equal the reciprocal of the *strength* of the underlying graph.

We efficiently compute a unique partition of the edges of the graph, called the prime-partition, and find the set of pure strategies of the hider that are best responses against every maxmin strategy of the wiretapper. Using these special pure strategies of the hider, which we call omni-connected-spanning-subgraphs, we define a partial order on the elements of the prime-partition. From the partial order, we obtain a linear number of simple two-variable inequalities that define the maxmin-polytope, and a characterization of its extreme points.

Our definition of the partial order allows us to find all equilibrium strategies of the wiretapper that minimize the number of pure best responses of the hider. Among these strategies, we efficiently compute the *unique* strategy that maximizes the least punishment that the hider incurs for playing a pure strategy that is not a best response. Finally, we show that this unique strategy is the nucleolus of the recently studied simple cooperative *spanning connectivity game*.

Keywords: cooperative game, network connectivity, network security, nucleolus, wiretapping, zero-sum game.

1 Introduction

Communication networks consist of two major layers, large static physical networks, and virtual networks built on top of them. The physical infrastructure comprises optical fibres, circuits, and routers etc., and rarely changes. A virtual network specifies how to route traffic between nodes, is software-driven and

^{*} This research was supported in part by EPSRC projects EP/D067170/1, EP/G064679/1, and by the Centre for Discrete Mathematics and its Applications (DIMAP), EPSRC award EP/D063191/1.

hence flexible. Modern physical networks are highly-connected and offer many possibilities for constructing virtual networks. Security is an important consideration for choosing a virtual network. One aspect of network security is resilience to wiretapping, which is the problem we study here from a game-theoretic perspective.

We consider the problem of maximizing the probability of hitting a strategically chosen hidden *virtual network* by placing a wiretap on a single link of a communication network, represented by an undirected, unweighted graph. This can be seen as a two-player win-lose (zero-sum) game that we call the *wiretap game*. A pure strategy of the wiretapper is an edge to tap, and of his opponent, the *hider*, a choice of virtual network, a *connected spanning subgraph*. The *wiretapper* wins, with payoff one, when he picks an edge in the network chosen by the hider, and loses, with payoff zero, otherwise. Thus, the *value* of this game is the greatest probability that the wiretapper can secure for hitting the hidden network. He does this by playing a maxmin strategy, which is a probability distribution on the edges. The value also equals the smallest probability that the hider can secure, which she does by playing a minmax strategy, which is a probability distribution on connected spanning subgraphs.

Our results. The value of the wiretap game is shown to equal the reciprocal of the *strength* of the underlying graph, a concept introduced by Gusfield [14]. We efficiently compute a unique partition of the edges of the graph, called the *prime-partition*. We find the set of pure strategies of the hider that are best responses against every maxmin strategy of the wiretapper. Using these special pure strategies of the hider, which we call *omni-connected-spanning-subgraphs*, we define a partial order on the elements of the prime-partition. Our definition in terms of omni-connected-spanning-subgraphs is central to proving our results.

From the partial order, we obtain a linear number of simple two-variable inequalities that define the maxmin-polytope, and a characterization of its extreme points. In contrast, the natural description of the maxmin-polytope is as the solutions to a linear program with exponentially many constraints. Our definition of the partial order allows us to find all equilibrium strategies of the wiretapper that minimize the number of pure best responses of the hider. Among these strategies, we efficiently compute the *unique* strategy that maximizes the least punishment that the hider incurs for playing a pure strategy that is not a best response.

Finally, we show that our analysis of the wiretap game provides a polynomial-time algorithm for computing the nucleolus of the *spanning connectivity game*, a simple cooperative game [6]. In this game, the players are the edges of the graph and a coalition, which is a subset of edges, has value one if it is a connected spanning subgraph, and zero otherwise. The characterization of the maxmin strategies of the wiretap game carries over to the least-core polytope of the spanning connectivity game, and the nucleolus of this game is the special maxmin strategy we compute for the wiretap game.

Related work. Wiretapping, as an important aspect of network security, has received recent attention in different settings, see e.g. [11] and [15].

The strength of an unweighted graph, which has a central role in our work, is also called the edge-toughness, and relates to the classical work of Nash-Williams [20] and Tutte [27]. Cunningham [5] generalized the concept of strength to edge-weighted graphs and proposed a strongly polynomial-time algorithm to compute it. Computing the strength of a graph is a special type of ratio optimization in the field of submodular function minimization [12]. Cunningham used the strength of a graph to address two different one-player optimization problems: the optimal attack and reinforcement of a network. The prime-partition we use is a truncated version of the principal-partition, first introduced by Narayanan [19] and Tomizawa [25]. The principal-partition was used in an extension of Cunningham’s work to an online setting [21]. Our work complements that of Cunningham and its successors by analyzing a new two-player game.

The nucleolus of the spanning connectivity game can be seen as a special maxmin strategy in the wiretap game. The connection between the nucleolus of a cooperative game and equilibrium strategies in a zero-sum game has been investigated before in a general context [22]. However, in many cases the nucleolus is hard to compute. The computational complexity of computing the nucleolus has attracted much attention [17], with both negative results [8, 10, 7], and positive results [13, 9, 16, 24]. Our positive results for the spanning connectivity game are in contrast to the negative results presented in [1], where it is shown that the problems of computing the Shapley values and Banzhaf values are $\#P$ -complete for the spanning connectivity game. Those results are a strengthening of the hardness results for the more general, min-base games, introduced in [18], and the positive results here thus apply to a special case of those games.

2 The wiretap game

The strategic form of the wiretap game is defined implicitly by the graph $G = (V, E)$. The pure strategies of the wiretapper are the edges E and the pure strategies of the hider are the set of connected spanning subgraphs \mathcal{S} . An element of \mathcal{S} is a set of edges, with a typical element denoted by S . The wiretapper receives payoff one if the edge he chooses is part of the spanning subgraph chosen by the hider, and receives payoff zero otherwise. Thus, the value of the game is the probability that the wiretapper can secure for wiretapping the connected spanning subgraph chosen by the hider.

Let $\Delta(A)$ be the set of mixed strategies (probability distributions) on a finite set A . By the well-known minmax theorem for finite zero-sum games, the wiretap game $\Gamma(G)$ has a unique *value*, defined by

$$val(\Gamma) = \max_{x \in \Delta(E)} \min_{S \in \mathcal{S}} \sum_{e \in S} x_e = \min_{y \in \Delta(\mathcal{S})} \max_{e \in E} \sum_{\{S \in \mathcal{S}: e \in S\}} y_S. \quad (1)$$

The equilibrium or *maxmin* strategies of the wiretapper are the solutions $\{x \in \Delta(E) \mid \sum_{e \in S} x_e \geq val(\Gamma) \text{ for all } S \in \mathcal{S}\}$ to the following linear program, which

has the optimal value $val(\Gamma)$.

$$\begin{aligned} & \max z \\ \text{s.t. } & \sum_{e \in S} x_e \geq z \text{ for all } S \in \mathcal{S}, \\ & x \in \Delta(E). \end{aligned} \quad (2)$$

Playing any maxmin strategy guarantees the wiretapper a probability of successful wiretapping of at least $val(\Gamma)$. The equilibrium or *minmax* strategies of the hider are $\{y \in \Delta(\mathcal{S}) \mid \sum_{\{S \in \mathcal{S}: e \in S\}} y_S \leq val(\Gamma) \text{ for all } e \in E\}$. Playing any minmax strategy guarantees the hider to suffer a probability of successful wiretapping of no more than $val(\Gamma)$. The following simple observation shows the importance of minimum connected spanning graphs in the analysis of the wiretap game. For a mixed strategy $x \in \Delta(E)$ and pure strategy $S \in \mathcal{S}$, the resulting probability of a successful wiretap is $\sum_{e \in S} x_e$. We denote by G^x the edge-weighted graph comprising the graph G with edge weights $x(e)$ for all $e \in E$. Let $w^*(x)$ be the weight of a minimum connected spanning graph of G^x .

Fact 1 *The set of pure best responses of the hider against the mixed strategy $x \in \Delta(E)$ is*

$$\{S \in \mathcal{S} \mid \sum_{e \in S} x_e = w^*(x)\}.$$

We could define the wiretap game by only allowing the hider to pick spanning trees, however, our definition with connected spanning subgraphs allows a clean connection to the spanning connectivity game.

3 Overview of results

In this section, we present our results. We start with the basic notations and definitions. From here on we fix a connected graph $G = (V, E)$. Unless mentioned explicitly otherwise, any implicit reference to a graph is to G and α is an edge-distribution, which is a probability distribution on the edges E . For ease, we often refer to the weighted graph G^α simply by α , where this usage is unambiguous. For a subgraph H of G , we denote by $\alpha(H)$ the sum $\sum_{e \in E(H)} \alpha(e)$, where $E(H)$ is the edge set of H . We refer to equilibrium strategies of the wiretapper as maxmin-edge-distributions.

Definition 1. *For every edge-distribution α , we denote its distinct weights by $x_1^\alpha > \dots > x_m^\alpha \geq 0$ and define $\mathcal{E}(\alpha) = \{E_1^\alpha, \dots, E_m^\alpha\}$ such that $E_i^\alpha = \{e \in E \mid \alpha(e) = x_i^\alpha\}$ for $i = 1, \dots, m$.*

Our initial goal is to characterize those partitions $\mathcal{E}(\alpha)$ that can arise from maxmin-edge-distributions α . We start with the following simple setting. Assume that the wiretapper is restricted to choosing a strategy α such that $|\mathcal{E}(\alpha)| = 2$, and $x_2^\alpha = 0$. Thus, the wiretapper's only freedom is the choice of the set E_1^α . What is his best possible choice? By Fact 1, a best response against α is a minimum connected spanning subgraph H of α . So the wiretapper should choose

E_1^α so as to maximize $\alpha(H)$. How can such an E_1^α be found? To answer, we relate the weight of a minimum connected spanning subgraph H of α to E_1^α .

To determine $\alpha(H)$, we may assume about H that for every connected component C of $(V, E \setminus E_1^\alpha)$ we have $E(H) \cap E(C) = E(C)$, since $\alpha(e) = 0$ for every $e \in E(C)$. We can also assume that $|E_1^\alpha \cap E(H)|$ is the number of connected components in $(V, E \setminus E_1^\alpha)$ minus 1, since this is the minimum number of edges in $E(H)$ that a connected spanning subgraph may have. To formalize this we use the following notation.

Definition 2. Let $E' \subseteq E$. We set $C_G(E')$, to be the number of connected components in the graph $G \setminus E'$, where $G \setminus E'$ is a shorthand for $(V, E \setminus E')$. If $E' = \emptyset$ we just write C_G .

Using the above notation, a connected spanning subgraph H is a minimum connected spanning subgraph of α if $|H \cap E_1^\alpha| = C_G(E_1^\alpha) - C_G = C_G(E_1^\alpha) - 1$. Now we can compute $\alpha(H)$. By definition, $x_1^\alpha = \frac{1}{|E_1^\alpha|}$ and $x_2^\alpha = 0$ and therefore

$$\alpha(H) = \frac{C_G(E_1^\alpha) - C_G}{|E_1^\alpha|}.$$

We call this ratio that determines $\alpha(H)$ the cut-rate of E_1^α . Note that it uniquely determines the weight of a minimum connected spanning subgraph of α .

Definition 3. Let $E' \subseteq E$. The cut-rate of E' in G is denoted by $cr_G(E')$ and defined as follows.

$$cr_G(E') := \begin{cases} \frac{C_G(E') - C_G}{|E'|} & \text{if } |V| > 1 \text{ and } |E'| > 0, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

We write $cr(E')$, unless we make a point of referring to a different graph.

Thus, when $|\mathcal{E}(\alpha)| = 2$ and $x_2^\alpha = 0$, a best choice of E_1^α is one for which $cr(E_1^\alpha)$ is maximum. Since E is finite, an E_1^α that maximizes $cr(E_1^\alpha)$ exists.

Definition 4. The cut-rate of G is defined as $opt := \max_{E' \subseteq E} cr(E')$.

By opt , we always refer to the cut-rate of the graph G . In case we refer to the cut-rate of some other graph, we add the name of the graph as a subscript. The value opt is a well known and studied attribute of a graph. It is equal to the reciprocal of the strength of a graph, as defined by Gusfield [14] and named by Cunningham [5]. There exists a combinatorial algorithm for computing the strength, and hence opt , that runs in time polynomial in the size of the graph, by which we always mean $|V| + |E|$.

We generalize the above technique to the case that α is not restricted. Assume again that H is a minimum connected spanning subgraph of α . Intuitively, even if α has more than 2 distinct weights we would expect $|E_1^\alpha \cap E(H)|$ to be as small as possible, i.e., $C_G(E_1^\alpha) - C_G$. We would also expect $|(E_1^\alpha \cup E_2^\alpha) \cap E(H)|$ to be as small as possible, i.e., $C_G(E_1^\alpha \cup E_2^\alpha) - C_G$. If these both hold then $|E_2^\alpha \cap E(H)| = C_G(E_1^\alpha \cup E_2^\alpha) - C_G(E_1^\alpha)$, which is the increase in the number of components we

get by removing the edges of E_2^α from $G \setminus E_1^\alpha$. Thus, the total weight contributed to H by edges in $E(H) \cap E(E_2^\alpha)$ is $x_2^\alpha (C_G(E_1^\alpha \cup E_2^\alpha) - C_G(E_1^\alpha))$. Now, unlike the previous case, we do not know x_2^α . However, this is not a problem since, as we shall see, we are interested in the ratio

$$\frac{\alpha(E(H) \cap E_2^\alpha)}{\alpha(E_2^\alpha)} = \frac{C_G(E_1^\alpha \cup E_2^\alpha) - C_G(E_1^\alpha)}{|E_2^\alpha|}.$$

We use the following notation to express this and its extension to more weights.

Definition 5. For $\ell = 1, \dots, |\mathcal{E}(\alpha)|$ we set

$$cr_\ell^\alpha = \frac{C_G(\cup_{i=1}^\ell E_i^\alpha) - C_G(\cup_{i=1}^{\ell-1} E_i^\alpha)}{|E_\ell^\alpha|}.$$

The intuition above indeed holds, as stated in the following proposition, which we prove in Appendix 0.

Proposition 1. Let H be a minimum connected spanning subgraph of α . Then $|E(H) \cap E_\ell^\alpha| = |E_\ell^\alpha| cr_\ell^\alpha$ for every ℓ such that $x_\ell^\alpha > 0$.

Using Proposition 1 we can relate the weight of a minimum connected spanning subgraph of α to the sets of $\mathcal{E}(\alpha)$. This relationship also characterizes the maxmin-edge-distributions, which are the edge-distributions whose minimum connected spanning subgraph weight is the maximum possible.

Theorem 1. Let H be a minimum connected spanning subgraph of α and $m = |\mathcal{E}(\alpha)|$. Then $\alpha(H) \leq opt$ and we have $\alpha(H) = opt$ if and only if

1. $cr_\ell^\alpha = opt$ for $\ell = 1, \dots, m-1$, and
2. if $cr_m^\alpha \neq opt$ then $x_m^\alpha = 0$.

Theorem 1 is proved in Appendix 1. An immediate implication of Theorem 1 is that opt is an upper bound on the value the wiretapper can achieve. This also follows from the well-known fact that the fractional packing number of spanning trees of a graph is equal to the strength of a graph, which in turn follows from the theorems of Nash-Williams [20] and Tutte [27] on the integral packing number (see also [3]). Since we have already seen that indeed the wiretapper can achieve opt by distributing all probability mass equally over an edge set that has cut-rate opt , we get the following.

Corollary 1. The value of the wiretap game is opt .

We know what the value of the game is and we know a characterization of the $\mathcal{E}(\alpha)$'s for maxmin-edge-distributions α . Yet this characterization does not give us a simple way to find maxmin-edge-distributions. Resolving this is our next goal. Since the set of maxmin-edge-distributions is convex, it is easy to show that there exists a maxmin-edge-distribution β such that for every $e_1, e_2 \in E$ we have $\beta(e_1) = \beta(e_2)$ if and only if $\gamma(e_1) = \gamma(e_2)$ for every maxmin-edge-distribution γ . This implies that $\mathcal{E}(\beta)$ refines $\mathcal{E}(\gamma)$ for every maxmin-edge-distribution γ , where by “refines” we mean the following.

Definition 6. Let $\mathcal{E}_1, \mathcal{E}_2$ be partitions of E . Then \mathcal{E}_1 refines \mathcal{E}_2 if for every set $E' \in \mathcal{E}_1$ there exists a set $E'' \in \mathcal{E}_2$ such that $E' \subseteq E''$.

Thus, there exists a partition of E that is equal to $\mathcal{E}(\beta)$ for some maxmin-edge-distribution β and refines $\mathcal{E}(\gamma)$ for every maxmin-edge-distribution γ . We call such a partition the *prime-partition*. It is unique since there can not be different partitions that refine each other.

Definition 7. The prime-partition \mathcal{P} is the unique partition that is equal to $\mathcal{E}(\beta)$ for some maxmin-edge-distribution β and refines $\mathcal{E}(\gamma)$ for every maxmin-edge-distribution γ .

Theorem 2. The prime-partition exists and can be computed in time polynomial in the size of G .

Theorem 2 is proved in Appendix 2. The prime-partition \mathcal{P} reveals a lot about the structure of the maxmin-edge-distributions. Yet by itself \mathcal{P} does not give us a simple means for generating maxmin-edge-distributions. Using the algorithm for finding \mathcal{P} one can show that, depending on G , there may be a unique element in \mathcal{P} whose edges are assigned 0 by every maxmin-edge-distribution.

Lemma 1. $cr_G(E) \neq opt$ if and only if there exists a unique set $D \in \mathcal{P}$ such that for every maxmin-edge-distribution α and $e \in D$ we have $\alpha(e) = 0$. If D exists then it can be found in time polynomial in the size of G .

Lemma 1 is proved in Appendix 3. From here on we shall always refer to the set D in Lemma 1 as the *degenerate set*. For convenience, if D does not exist then we shall treat both $\{D\}$ and D as the empty set. See Figure 1 for an example of the prime-partition and the degenerate set.

We use the prime-partition to define a special subset of the minimum connected spanning subgraphs that we call the *omni-connected-spanning-subgraphs*, which are useful for proving the characterization of maxmin-edge-distributions and their refinements.

Definition 8. A connected spanning subgraph H is an omni-connected-spanning-subgraph if for every $P \in \mathcal{P} \setminus \{D\}$ we have

$$|E(H) \cap P| = |P| \cdot opt .$$

Proposition 2. There exists an omni-connected-spanning-subgraph.

Proof. Let β be a maxmin-edge-distribution such that $\mathcal{E}(\beta) = \mathcal{P}$. Let H be a minimum connected spanning subgraph of β . Then by Proposition 1, we have that H is an omni-connected-spanning-subgraph. \square

The omni-connected-spanning-subgraphs are the set of the hider's pure strategies that are best responses against every maxmin-edge-distribution.

Proposition 3. For every edge-distribution α such that \mathcal{P} refines $\mathcal{E}(\alpha)$ and $\alpha(e) = 0$ for every $e \in D$ and omni-connected-spanning-subgraph H , we have $\alpha(H) = opt$.

We prove Proposition 3 in Appendix 4. The importance of omni-connected-spanning-subgraphs stems from the following scenario. Assume that \mathcal{P} refines $\mathcal{E}(\alpha)$ and $\alpha(e) = 0$ for every $e \in D$, and let H be an omni-connected-spanning-subgraph. By Proposition 3, we know that $\alpha(H) = \text{opt}$. Suppose we can remove from H an edge from $E(H) \cap P$, where P is a nondegenerate element of \mathcal{P} , and add a new edge from another set $P' \setminus E(H)$ in order to get a new connected spanning subgraph. Assume α assigns to the edge removed strictly more weight than it assigns to the edge added. Then the new connected spanning subgraph has weight strictly less than $\alpha(H)$ and hence strictly less than opt , since $\alpha(H) = \text{opt}$ by Proposition 3. Consequently, α is not a maxmin-edge-distribution and we can conclude that any edge-distribution β that assigns to each edge in P strictly more weight than to the edges in P' is not a maxmin-edge-distribution. This intuition is captured by the following definition, which leads to the characterization of maxmin-edge-distributions in Theorem 3.

Definition 9. *Let $P, P' \in \mathcal{P} \setminus \{D\}$ be distinct. Then P leads to P' if and only if there exists an omni-connected-spanning-subgraph H with $e \in P \setminus E(H)$ and $e' \in P' \cap E(H)$ such that $(H \setminus \{e'\}) \cup \{e\}$ is a connected spanning subgraph. We denote the “leads to” relation by \mathcal{R} .*

Definition 10. *An edge-distribution α agrees with \mathcal{R} if \mathcal{P} refines $\mathcal{E}(\alpha)$ and for every $P \in \mathcal{P} \setminus \{D\}$ that is a parent of $P' \in \mathcal{P} \setminus \{D\}$ and $e \in P$, $e' \in P'$ we have $\alpha(e) \geq \alpha(e')$, and for every $e \in D$ we have $\alpha(e) = 0$.*

Theorem 3. *An edge-distribution α is a maxmin-edge-distribution if and only if it agrees with \mathcal{R} .*

Theorem 3 is proved in Appendix 5.1. By definition, there exists a maxmin-edge-distribution β with $\mathcal{E}(\beta) = \mathcal{P}$. By Theorem 3, we have that β agrees with \mathcal{R} and hence the following holds.

Proposition 4. *The relation \mathcal{R} is acyclic.*

This allows us to define the acyclic parent-child relation, which is a simplification of \mathcal{R} and easy to find.

Definition 11. *Let $P, P' \in \mathcal{P} \setminus \{D\}$ be distinct. We say that P is a parent of P' (conversely P' a child of P) if P leads to P' and there is no $P'' \in \mathcal{P}$ such that P leads to P'' and P'' leads to P' . We refer to the relation as the parent-child relation and denote it by \mathcal{O} .*

The following is an immediate corollary of Theorem 3 and Definition 11.

Corollary 2. *An edge-distribution α is a maxmin-edge-distribution if and only if it agrees with \mathcal{O} .*

See Figure 1 for an example of an omni-connected-spanning-subgraph and the exchangeability of edges between a parent and child. Corollary 2 defines a linear inequality for each parent and child in the relation \mathcal{O} . Along with the inequalities

that define a probability distribution on edges, this gives a small number of two-variable inequalities describing the maxmin-polytope. In Appendix 8 we characterize the extreme points of maxmin-polytope. The proof of the following theorem, which states that \mathcal{O} can be found in polynomial time, can be found in Appendix 5.2.

Theorem 4. *The parent-child relation \mathcal{O} can be computed in time polynomial in the size of G .*

The wiretapper will in general have a choice of infinitely many maxmin-edge-distributions. To choose a maxmin-edge-distribution, it is natural to consider refinements of the Nash equilibrium property that are beneficial to the wiretapper if the hider does not play optimally. First we show how to minimize the number of pure best responses of the hider. To do this, we use the relation \mathcal{O} to characterize a special type of maxmin-edge-distribution which achieves this. We call this a prime-edge-distribution. The prime-edge-distributions are characterized by the following lemma.

Definition 12. *A maxmin-edge-distribution α is a prime-edge-distribution if the number of the hider's pure best responses against it is the minimum possible.*

Lemma 2. *An edge-distribution γ is a prime-edge-distribution if and only if $\gamma(e) > 0$ for every $e \in E \setminus D$, and for every $P, P' \in \mathcal{P} \setminus \{D\}$ such that P is a parent of P' and every $e \in P, e' \in P'$, we have $\gamma(e') > \gamma(e'')$.*

Using this characterization one can easily check whether α is a prime-edge-distribution and one can also easily construct a prime-edge-distribution.

We prove Lemma 2 in Appendix 6. The proof runs as follows. First we show that for any α that satisfies the condition of the lemma, every minimum connected spanning subgraph is an omni-connected-spanning-subgraph. Hence, using Proposition 3, we get that for any α that satisfies the condition of the lemma, a connected spanning subgraph is a minimum connected spanning subgraph of α if and only if it is an omni-connected-spanning-subgraph. These are the only such maxmin-edge-distributions, since any maxmin-edge-distribution that does not satisfy the condition of the lemma has a parent and its child whose edges get the same weight. Consequently, by the definition of parent and child, it has a minimum connected spanning subgraph that is not an omni-connected-spanning-subgraph.

We have already seen how to minimize the number of pure best responses of the hider, by playing a prime-edge-distribution. We now show how to uniquely maximize the weight of a pure second-best response by choosing between prime-edge-distributions. This maximizes the least punishment that the hider will incur for picking a non-optimal pure strategy.

Against a prime-edge-distribution, the candidates for pure second-best responses are those connected spanning subgraphs that differ from omni-connected-spanning-subgraphs in at most two edges. For each parent and child we have at least one of these second-best responses. A second-best response either is a best




Fig. 1. LEFT: The left figure illustrates the prime-partition $\mathcal{P} = \{E_1, \dots, E_5\}$. For this graph, $\text{opt} = 1/2$. The set $E_1 = \{e_1, e_2\}$, the set $E_2 = \{e_3, e_4\}$, the set E_3 is equal to the edges of the left K_4 , the set E_4 is equal to the edges of the right K_4 , and the set E_5 is equal to the edges of the K_5 . Suppose that maximin-edge-distribution β is such that $\mathcal{E}(\beta) = \mathcal{P}$, and $E_i^\beta = E_i$ for $i = 1, \dots, 5$. (There will be other maximin-edge-distributions with the same partition in which E_3 and E_4 exchange roles.) Removing E_1 from the graph creates one extra component by removing two edges, so we have $\text{cr}_1^\beta = \text{cr}(E_1) = \text{opt} = 1/2$. Similarly we have $\text{cr}_k^\beta = 1/2$ for all $k = 1, \dots, 4$. However, $\text{cr}_5^\beta = 4/10 < 1/2$ and so the set E_5 is a degenerate set, as per Lemma 1. The figure shows the subgraph H indicated with solid edges. It is an omni-connected-spanning-subgraph, using two edges from each of the K_4 's, one edge from the two edges that connect the two K_4 's, and one edge from the two edges that connect the two K_4 's to the K_5 . Within the K_5 , an omni-connected-spanning-subgraph can use more than four edges, as this K_5 corresponds to the final element of the prime-partition with any strong linear order and achieves cut-rate $4/10$, which is worse than $\text{opt} = 1/2$. The edge e_3 can be replaced with the edge e_1 . Thus, the edges in the element of the prime-partition containing e_1 must have weight at least that of the edges in the element of the prime-partition containing e_3 . RIGHT: The right figure illustrates the partial order \mathcal{O} and its layers $\{L_1, L_2, L_3\}$.

response with one extra edge, or it differs from a best response in two edges, where it has one less edge in a child of \mathcal{O} and one more in the child's parent.

We are only interested in the case that $\text{opt} < 1$, since the graph has $\text{opt} = 1$ if and only if it contains a bridge, in which case the value of the game is one and the hider does not have a second-best response. From here on we assume the following.

Assumption 1 $\text{opt} < 1$.

Intuitively, to maximize the weight of a second-best response, we want to minimize the number of distinct weights. The minimum number of distinct positive weights we can achieve for a prime-edge-distribution is equal to the number of elements in the longest chain in the parent-child relation. This motivates the following definition.

Definition 13. We define $\mathcal{L}_1, \mathcal{L}_2, \dots$ inductively as follows. The set \mathcal{L}_1 is all the sinks of \mathcal{O} excluding D . For $j = 2, \dots$, we have that \mathcal{L}_j is the set of all the sinks when all elements of $\{D\} \cup (\cup_{i=1, \dots, j-1} \mathcal{L}_i)$ have been removed from \mathcal{O} .

Note that \mathcal{O} is defined only over nondegenerate elements of \mathcal{P} and hence the degenerate set is not contained in any of $\mathcal{L}_1, \mathcal{L}_2, \dots$.

Definition 14. The layers $\mathcal{L} = \{L_1, \dots, L_t\}$ of G are $L_i = \cup_{E' \in \mathcal{L}_i} E'$ for $i = 1, \dots, t$.

See Figure 1 for an example of layers. The following theorem shows that there is a unique maxmin-edge-distribution that maximizes the difference between the payoff of a best and second-best response. This unique maxmin-edge-distribution turns out to be the nucleolus of the spanning connectivity game, as explained in Section 4. For convenience, we refer to this strategy as the nucleolus.

Theorem 5. Let

$$\kappa = \frac{1}{\sum_{i=1}^t i \cdot |L_i|} .$$

The nucleolus ν has $\nu(e) = i \cdot \kappa$ for every $i \in \{1, \dots, t\}$ and $e \in L_i$ and $\nu(e) = 0$ otherwise.

Theorem 5 is proved in Appendix 7. The proof says that the weight of a second-best response is $\text{opt} + \kappa$, and this must be optimal since all the weights are multiples of κ . For all other prime-edge-distributions there is a second-best response with weight less than $\text{opt} + \kappa$.

4 Spanning connectivity games

A simple cooperative game (N, v) consists of a player set $N = \{1, \dots, n\}$ and characteristic function $v : 2^N \rightarrow \{0, 1\}$ with $v(\emptyset) = 0$, $v(N) = 1$, and $v(S) \leq v(T)$ whenever $S \subseteq T$. A coalition $S \subseteq N$ is *winning* if $v(S) = 1$ and *losing* if $v(S) = 0$. The payoff vector to the players $x = (x_1, \dots, x_n)$ satisfies $x(N) = v(N) = 1$, with $x(S) = \sum_{i \in S} x_i$. It is called an *imputation* if $x_i \geq v(\{i\})$ for all $i \in N$. For a game (N, v) and imputation $x = (x_1, \dots, x_n)$, the *excess* $e(x, S)$ of a coalition S under x is defined as $e(x, S) = x(S) - v(S)$.

We relate our analysis of the wiretap game to two cooperative solutions based on the excess of coalitions: the *least core* and the *nucleolus*, which is a unique point in the least core. An imputation x is in the ϵ -core if $e(x, S) \geq -\epsilon$ for all $S \subset N$. An imputation x is in the *least core* if it is in the ϵ -core for the smallest possible ϵ . The *excess vector* of an imputation x is $(e(x, S_1), \dots, e(x, S_{2^n}))$, where $e(x, S_1) \leq e(x, S_2) \leq \dots \leq e(x, S_{2^n})$. The *nucleolus* is the element of the least core which has the largest excess vector lexicographically. The nucleolus is unique [23]. We denote the *distinct* excesses by $(-\epsilon_1, \dots, -\epsilon_t)$ for $t \leq 2^n$, where $\epsilon_1 > \dots > \epsilon_t$. Note that, by definition, $\epsilon_1 = \epsilon$, the least core value, which is the optimal value of the objective function in (4).

For a graph $G = (V, E)$ with at least three nodes, the *spanning connectivity game*, introduced in [1], has player set E and characteristic function

$$v(S) = \begin{cases} 1, & \text{if there exists a spanning tree } T = (N, E') \text{ such that } E' \subseteq S, \\ 0, & \text{otherwise.} \end{cases}$$

Since the graph has at least three nodes, $v(\{i\}) = 0$ for all $i \in E$, so x is an imputation if and only if it is a probability distribution on players, i.e., $x \in \Delta(E)$. The least core of the spanning connectivity game is the set of all solutions to the following linear program:

$$\begin{aligned} \min \quad & \epsilon \\ \text{s.t.} \quad & e(x, S) \geq -\epsilon, \text{ for all } S \subset E, \\ & x \in \Delta(E). \end{aligned} \tag{4}$$

First we show that the least core is identical to the maxmin-polytope.

Proposition 5. *The least core of the spanning connectivity game is the set of maxmin-edge-distributions of the wiretap game.*

Proof. The problems of finding a maxmin-edge-distribution and an element of the least core are given by the LPs (2) (from Section 2) and (4), respectively. The solution of (4) satisfies $\epsilon \geq 0$, since we have $x(S) - v(S) \geq -\epsilon$ and $x(S) \leq 1$, and for any winning coalition $v(S) = 1$. So, for any losing coalition S , where $v(S) = 0$, the inequality $e(x, S) = x(S) - v(S) \geq -\epsilon$ in (4) is redundant, and only the inequalities for winning coalitions, i.e., connected spanning subgraphs are needed. Note that $x(S) = \sum_{e \in S} x_e$. Hence, the linear programs (2) and (4) have the same solutions with $z = 1 - \epsilon$, except for the objective functions that differ only by a constant. \square

Now we show that the nucleolus is the unique maxmin-edge-distribution that minimizes the number of pure best responses of the hider and, given this, maximizes the probability arising from the hider playing a pure second-best response.

Proposition 6. *The nucleolus of the spanning connectivity game is identical to the maxmin-edge-distribution defined in Theorem 5.*

Proof. The maxmin-edge-distribution ν in Theorem 5 minimizes the number of pure best responses of the hider, i.e., it minimizes the number of minimum connected spanning subgraphs in G^ν . This is equivalent to minimizing the number of ϵ_1 -coalitions in the spanning connectivity game. Moreover, it *uniquely* maximizes the probability for a successful wiretap of a second-best response of the hider, which is equivalent to maximizing ϵ_2 . \square

5 Further research

The equilibrium strategies of the hider can be found using the complete refined principal partition [2] (our prime-partition is a truncation of this one). A characterization of these minmax strategies would be an interesting next step.

Extensions to wiretap game. There are a number of natural extensions to the wiretap game. For example, if the wiretapper is allowed to pick multiple edges. In Figure 1, if the wiretapper can pick two edges, then by choosing e_1 and e_2 , he guarantees success. With the number of edges to pick as input, is this problem computationally tractable, or hard? One could consider variants where the nodes of the hider are a subset of all nodes, unknown to the wiretapper. Another natural extension is to make the game dynamic with multiple rounds.

Further equilibrium refinements. Potters and Tijs [22] define the “nucleolus of a matrix game” and show that for a matrix the nucleolus, which is no longer unique as for a cooperative game, corresponds to the *proper equilibria* of the matrix game. This equilibrium concept, unlike Nash equilibria for zero-sum games, is not independent for the two players (for a Nash equilibrium, one player can independently play any maxmin strategy and the other any minmax strategy in equilibrium). Unlike the special strategy of the wiretapper we compute here, computing a proper equilibrium will require a simultaneous analysis of the strategies of both the wiretapper and hider, however it seems plausible that the structure we have shown here may be enough to do this. So, can we efficiently find one proper equilibrium of the wiretap game, or even a characterization of the complete set of proper equilibria? What about other equilibrium refinements?

References

1. H. Aziz, O. Lachish, M. Paterson, and R. Savani. Power indices in spanning connectivity games. In *AAIM: Algorithmic Aspects in Information and Management*, pages 55–67, 2009.
2. P. A. Catlin, J. W. Grossman, A. M. Hobbs, and H.-J. Lai. Fractional arboricity, strength, and principal partitions in graphs and matroids. *Discrete Appl. Math.*, 40(3):285–302, 1992.
3. D. Chakrabarty, A. Mehta, and V. V. Vazirani. Design is as easy as optimization. In *ICALP: International Colloquium on Automata, Languages and Programming*, pages 477–488, 2006.
4. E. Cheng and W. H. Cunningham. A faster algorithm for computing the strength of a network. *Inf. Process. Lett.*, 49(4):209–212, 1994.
5. W. H. Cunningham. Optimal attack and reinforcement of a network. *J. ACM*, 32(3):549–561, 1985.
6. X. Deng and Q. Fang. Algorithmic cooperative game theory. In *Pareto Optimality, Game Theory And Equilibria*, volume 17 of *Springer Optimization and Its Applications*, pages 159–185. Springer, 2008.
7. X. Deng, Q. Fang, and X. Sun. Finding nucleolus of flow game. In *SODA '06: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 124–131, New York, NY, USA, 2006. ACM.
8. E. Elkind, L. Goldberg, P. Goldberg, and M. Wooldridge. Computational complexity of weighted threshold games. *AAAI-07 (Twenty-Second National Conference on Artificial Intelligence)*, 2007.
9. E. Elkind and D. Pasechnik. Computing the nucleolus of weighted voting games. In *SODA '09: Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 327–335, Philadelphia, PA, USA, 2009.
10. U. Faigle, W. Kern, and J. Kuipers. Computing the nucleolus of min-cost spanning tree games is NP-hard. *Int. J. Game Theory*, 27(3):443–450, 1998.
11. M. K. Franklin, Z. Galil, and M. Yung. Eavesdropping games: a graph-theoretic approach to privacy in distributed systems. *J. ACM*, 47(2):225–243, 2000.
12. S. Fujishige. *Submodular Functions and Optimization*, volume 58 of *Annals of Discrete Mathematics*. Elsevier, 2005.
13. D. Granot, M. Maschler, G. Owen, and W. R. Zhu. The kernel/nucleolus of a standard tree game. *International Journal of Game Theory*, 25(2):219–44, 1996.
14. D. Gusfield. Connectivity and edge-disjoint spanning trees. *Inf. Process. Lett.*, 16(2):87–89, 1983.
15. K. Jain. Security based on network topology against the wiretapping attack. *IEEE Wireless Communications*, 11(1):68 – 71, 2004.
16. J. Kuipers. A polynomial algorithm for computing the nucleolus of convex games. *Tech. rep., Univ. of Maastricht*, M 96-12, July 1996.
17. J. Kuipers, U. Faigle, and W. Kern. On the computation of the nucleolus of a cooperative game. *International Journal of Game Theory*, 30(1):79–98, 2001.
18. H. Nagamochi, D.-Z. Zeng, N. Kabutoya, and T. Ibaraki. Complexity of the minimum base game on matroids. *Math. Oper. Res.*, 22(1):146–164, 1997.
19. H. Narayanan. *Theory of matroids and network analysis*. PhD thesis, IIT, Bombay, 1974.
20. C. S. A. Nash-Williams. Edge-disjoint spanning trees of finite graphs. *J. London Math. Soc.*, 36:445 – 450, 1961.
21. S. Patkar and H. Narayanan. Fast on-line/off-line algorithms for optimal reinforcement of a network and its connections with principal partition. In *FST TCS: Foundations of Software Technology and Theoretical Computer Science*, pages 94–105. Springer-Verlag, 2000.

22. J. A. M. Potters and S. H. Tijs. The nucleolus of a matrix game and other nucleoli. *Math. Oper. Res.*, 17(1):164–174, 1992.
23. D. Schmeidler. The nucleolus of a characteristic function game. *SIAM J. Appl. Math.*, 17(6):1163–1170, 1969.
24. T. Solymosi and T. E. S. Raghavan. An algorithm for finding the nucleolus of assignment games. *Int. J. Game Theory*, 23(2):119–143, 1994.
25. N. Tomizawa. Strongly irreducible matroids and principal partition of a matroid into strongly irreducible minors. *Electron. and Commun. Japan*, 59(A):1–10, 1976.
26. V. A. Trubin. Strength of a graph and packing of trees and branchings. *Cybernetics and Systems Analysis*, 29(3):379–384, 1993.
27. W. T. Tutte. On the problem of decomposing a graph into n connected factors. *J. London Math. Soc.*, 36:221 – 230, 1961.

0 Appendix: Preliminaries

Proof of Proposition 1

Let H be a minimum connected spanning subgraph of α . And let t be the maximum such that $x_t^\alpha > 0$. We next show that $|E(H) \cap E_i^\alpha| = |E_i^\alpha|cr_i^\alpha$ for $i = 1, \dots, t$.

Assume for the sake of contradiction that this is not so. Let k be minimal such that $|E(H) \cap E_k^\alpha| \neq |E_k^\alpha|cr_k^\alpha$. By the minimality of k we have

$$|E(H) \cap (\cup_{i=1}^{k-1} E_i^\alpha)| = \sum_{i=1}^{k-1} |E_i^\alpha|cr_i^\alpha. \quad (5)$$

Set $E' = \cup_{i=1}^k E_i^\alpha$. By the definition of cut-rate the number of connected components in $G \setminus E'$ is

$$C_G(E') = 1 + \sum_{i=1}^k |E_i^\alpha|cr_i^\alpha. \quad (6)$$

Thus $|E(H) \cap E'|$ is at least $\sum_{i=1}^k |E_i^\alpha|cr_i^\alpha$ and therefore by (5) we have $|E(H) \cap E_k^\alpha| \geq |E_k^\alpha|cr_k^\alpha$.

Assume $|E(H) \cap E_k^\alpha| > |E_k^\alpha|cr_k^\alpha$. Then, by (5), we have

$$|E(H) \cap E'| > \sum_{i=1}^k |E_i^\alpha|cr_i^\alpha. \quad (7)$$

We show next that this implies that there exists a connected spanning subgraph whose weight by α is strictly less than $\alpha(H)$ in contradiction to H being a minimum connected spanning subgraph. Set $s = C_G(E')$ and let C_1, \dots, C_s be the connected components of $G \setminus E'$. Now as H is a minimum connected spanning subgraph the set of edges in $E(H) \cap E'$ does not have a cycle, otherwise we could have removed one of them to get a connected spanning subgraph with strictly less weight. Thus the number of connected components of $E(H) \setminus E'$ is $1 + |E \cap E(H)|$. Set $r = |E' \cap E(H)|$ and let H_1, \dots, H_r be the connected components of $H \setminus E'$.

Note that for each $i \in \{1, \dots, r\}$ there exists a unique $j \in \{1, \dots, s\}$ such that $E(H_i) \subseteq E(C_j)$. For each $j \in \{1, \dots, s\}$ set I_j to be the set of all $i \in \{1, \dots, r\}$ such that $E(H_i) \subseteq E(C_j)$. By (6) and (7) we have $s < r$ and therefore by the pigeon-hole principle there exists $j \in \{1, \dots, r\}$ such that $|I_j| > 1$. Since C_j is a connected component and H a connected spanning subgraph there exist $x, y \in I_j$ and $e = \{u, v\} \in E(C_j) \setminus \cup_{i=1}^{|I_j|} E(H_i)$ such that $u \in V(H_x)$ and $v \in V(H_y)$. Again because H is a connected spanning subgraph there is a path in H between u and v this path contains edges not in $E(C_j)$ because u, v are in different connected components of $H \setminus E'$. Thus this path contains an edge $e' \in E'$ because only edges from E' connect the vertices of C_j to the rest of the graph. Consequently $(H \setminus \{e'\}) \cup \{e\}$ is a connected spanning subgraph. Since $e \notin E'$ we have $\alpha(e) < \alpha(e')$ and consequently $\alpha(H) > \alpha((H \setminus \{e'\}) \cup \{e\})$.

Fact 2 *Let H be a minimum connected spanning subgraph of α and $m = \mathcal{E}(\alpha)$ then $\alpha(H) = \sum_{i=1}^m x_i^\alpha |E_i^\alpha|cr_i^\alpha$ and for each $i = 1, \dots, m$ if $cr_i^\alpha < 1$ then there exists $e \in E_i^\alpha \setminus E(H)$.*

Proof. By definition $|E(H)| = \sum_{i=1}^m |E(H) \cap E_i^\alpha|$. Therefore $\alpha(H) = \sum_{i=1}^m x_i^\alpha |E(H) \cap E_i^\alpha|$. By applying Proposition 1 we get that $\alpha(H) = \sum_{i=1}^m x_i^\alpha |E_i^\alpha|cr_i^\alpha$.

Fix $i \in \{1, \dots, m\}$. By Proposition 1 we have $|E(H) \cap E_i^\alpha| = |E_i^\alpha|cr_i^\alpha$ and hence if $cr_i^\alpha < 1$ then $|E(H) \cap E_i^\alpha| < |E_i^\alpha|$ and therefore $E_i^\alpha \setminus E(H)$ is not empty. \square

Fact 3 Let $E_1, \dots, E_s \subseteq E$ be such that $E_i \cap E_j = \emptyset$ for every distinct $i, j \in \{1, \dots, s\}$. For $\ell = 1, \dots, s$ let r_ℓ be the cut-rate of E_ℓ in $G \setminus \bigcup_{i=1}^{\ell-1} E_i$. Assume that $r_\ell \geq y$ ($r_\ell \leq y$) for each $\ell = 1, \dots, s$. Then if there exists $i \in \{1, \dots, s\}$ such that $r_i > y$ ($r_i < y$) we have $cr(\bigcup_{i=1}^s E_i) > y$ ($r < y$) and otherwise $cr(\bigcup_{i=1}^s E_i) = y$.

Proof. By the definition of cut-rate $C_G(\bigcup_{i=1}^s E_i) = C_G + \sum_{i=1}^s r_i |E_i|$ and hence

$$cr(\bigcup_{i=1}^s E_i) = \frac{(C_G + \sum_{i=1}^s x_i |E_i|) - C_G}{\sum_{j=1}^s |E_j|} \geq \frac{\sum_{i=1}^s y |E_i|}{\sum_{j=1}^s |E_j|} = y.$$

Note that the above inequality is strict unless $r_i = y$ for $i = 1, \dots, s$. The proof for the case that $r_i \leq y$ for $i = 1, \dots, s$, is the same. \square

Definition 15. A minimal set $E' \subseteq E$ such that $cr(E') = opt$ is a prime-set of G .

Proposition 7. For $E', E'' \subset E$ such that $cr(E') = cr(E'') = opt$ the following holds:

1. $opt_{G \setminus E'} \leq opt$.
2. If $E'' \neq E'$ then $cr_{G \setminus E'}(E'' \setminus E') = opt$.
3. If $E'' \cap E' \neq \emptyset$ then $cr(E'' \cap E') = opt$.
4. If $E'' \setminus E' \neq \emptyset$ then $opt_{G \setminus E'} = opt$.
5. If E' is a prime-set then either $E' \subseteq E''$ or $E' \cap E'' = \emptyset$.

Proof. Note that $opt = 0$ only if $E = \emptyset$ and therefore in this case the proposition trivially holds. Assume that $opt > 0$. Hence by the definition of cut-rate we have $E', E'' \neq \emptyset$. We shall also assume that $E' \neq E''$ since otherwise the last four items hold trivially. We next prove the first item.

Let $E^* \subseteq E \setminus E'$ be such that $cr_{G \setminus E'}(E^*) = opt_{G \setminus E'}$. By definition such a set exists. Observe that $cr_{G \setminus E'}(E^*) \leq opt$ because otherwise since $cr(E') = opt$ by Fact 3, we have $cr(E' \cup E^*) > opt$, which is a contradiction to the maximality of opt .

We now prove the second and third items. If $E'' \cap E' = \emptyset$ then both items trivially hold. Assume $E'' \cap E' \neq \emptyset$. According to the first item $cr_{G \setminus E'}(E'' \setminus E') \leq opt$ and by definition $cr(E'' \cap E') \leq opt$. Thus as $cr(E'') = opt$ by Fact 3 we have $cr_{G \setminus E'}(E'' \setminus E') = cr(E'' \cap E') = opt$.

Finally we prove the last two items. Assume $E'' \setminus E' \neq \emptyset$. By the first item $opt_{G \setminus E'} \leq opt$. By the second item $cr_{G \setminus E'}(E'' \setminus E') = opt$ and hence also $opt_{G \setminus E'} \geq opt$ and consequently $opt_{G \setminus E'} = opt$.

Assume that E' is a prime-set. If $E' \cap E'' \neq \emptyset$ then by the second item $cr(E' \cap E'') = opt$ and hence by the definition of prime-set $E' \cap E'' = E'$ which implies $E' \subseteq E''$. \square

1 Appendix: Proof of Theorem 1

Let β be a edge-distribution and $s = |\mathcal{E}(\beta)|$. We say β is *strong* if $cr_\ell^\beta = opt$ for $\ell = 1, \dots, s-1$ and if $cr_s^\beta \neq opt$ then $x_s^\beta = 0$. From here on in this section H is a minimum connected spanning subgraph of α . Assume α is strong. By Fact 2

$$\alpha(H) = \sum_{\ell=1}^{|\mathcal{E}(\alpha)|} x_\ell^\alpha |E_\ell^\alpha| cr_\ell^\alpha.$$

Therefore as we have $cr_i^\alpha = opt$ for every i such that $x_i^\alpha > 0$ we conclude

$$\alpha(H) = opt \sum_{\ell=1}^{|\mathcal{E}(\alpha)|} x_\ell^\alpha |E_\ell^\alpha|.$$

Finally, since α is an edge-distribution $\sum_{\ell=1}^{|\mathcal{E}(\alpha)|} x_\ell^\alpha |E_\ell^\alpha| = 1$, we get that $\alpha(H) = \text{opt}$. Now the theorem directly follows from the subsequent lemma. \square

Lemma 3. *Let H be a minimum connected spanning subgraph of α , then $\alpha(H) \leq \text{opt}$ and if $\alpha(H) = \text{opt}$ then α is strong.*

Intuition for the proof Lemma 3. The proof of the Lemma 3 is by induction on s , the maximum index such that $x_s^\alpha > 0$. The basis of the induction is straightforward. The induction assumption states that for an edge-distribution β with $s-1$ distinct strictly positive weights, and minimum connected spanning subgraph H' of β we have $\beta(H') \leq \text{opt}$ and if $\beta(H') = \text{opt}$ then β is strong.

The main idea in the induction step is to show that one can shift around some of the weight of α in order to get a new edge-distribution β , such that β has exactly $s-1$ strictly positive distinct weights and $\beta(H') \geq \alpha(H)$ (or $\beta(H') > \alpha(H)$), where H' is a minimum connected spanning subgraph of β . Now since β has $s-1$ strictly positive distinct weights the induction assumption applies to it and hence $\beta(H) \leq \text{opt}$. This in turn implies that $\alpha(H) \leq \text{opt}$. Now by the above if $\alpha(H) = \text{opt}$ then also $\beta(H) = \text{opt}$ and hence by the induction assumption β is strong. With a bit of extra work this leads to α being strong.

The induction step consists of two separate cases. In the first case it is assumed that $cr_s^\alpha \leq cr(\cup_{i=1}^{s-1} E_i^\alpha)$, in the second $cr_s^\alpha > cr(\cup_{i=1}^{s-1} E_i^\alpha)$.

In the first case, by taking all the weight assigned by α to the edges of E_s^α and distributing it equally among the edges in $\cup_{i=1}^{s-1} E_i^\alpha$, one gets a new edge-distribution γ that has $s-1$ distinct strictly positive weights and $\alpha(H) \leq \gamma(H')$, where H' is a minimum connected spanning subgraph of γ . In the second case, one obtains the new edge-distribution from α in the following way. A constant amount of weight χ is reduced from each edge in $\cup_{i=1}^{s-1} E_i^\alpha$ and divide the total removed weight $\chi |\cup_{i=1}^{s-1} E_i^\alpha|$ equally among the edges of E_s^α thus getting a new edge-distribution δ where $\alpha(H) < \delta(H'')$, where H'' is a minimum connected spanning subgraph of δ . The value of χ is chosen so that δ gives the same weight to all the edges in $E_s^\alpha \cup E_{s-1}^\alpha$. Therefore the number of strictly positive weights of δ is $s-1$.

Proof of Lemma 3. If $s = 1$ then by Proposition 1 we have $\alpha(H) = x_1^\alpha |E_1^\alpha| cr_1^\alpha = cr_1^\alpha = cr(E_1^\alpha) \leq \text{opt}$. Note that if equality holds then $cr_1^\alpha = \text{opt}$ and hence α is strong.

Let $s > 1$. The induction assumption is that for every edge-distribution β that has $s-1$ strictly positive weights, we have $\beta(H') \leq \text{opt}$ for H' that is a minimum connected spanning subgraph of β and if $\beta(H') = \text{opt}$ then β is strong.

For the inductive step we deal with two cases separately. In case **(a)** we assume that

$$cr_s^\alpha \leq \frac{\sum_{i=1}^{s-1} cr_i^\alpha |E_i^\alpha|}{\sum_{i=1}^{s-1} |E_i^\alpha|}. \quad (8)$$

In case **(b)** we assume that

$$cr_s^\alpha > \frac{\sum_{i=1}^{s-1} cr_i^\alpha |E_i^\alpha|}{\sum_{i=1}^{s-1} |E_i^\alpha|}. \quad (9)$$

Note that we chose to write the more cumbersome $\frac{\sum_{i=1}^{s-1} cr_i^\alpha |E_i^\alpha|}{\sum_{i=1}^{s-1} |E_i^\alpha|}$ instead of $cr(\cup_{i=1}^{s-1} E_i^\alpha)$ as this form serves our purpose better.

(a) Set

$$\rho = \frac{x_s^\alpha |E_s^\alpha|}{\sum_{i=1}^{s-1} |E_i^\alpha|},$$

which is the total weight of E_s^α divided equally among all edges in $\cup_{i=1}^{s-1} E_i^\alpha$. Define $\gamma : E(G) \rightarrow \mathbb{R}$ so that $\gamma(e) = \alpha(e) + \rho$ for every $e \in \cup_{j=1}^{s-1} E_j^\alpha$ and $\gamma(e) = 0$ for every $e \in E(G) \setminus \cup_{j=1}^{s-1} E_j^\alpha$. According to this definition

$$\sum_{e \in E} \gamma(e) = \sum_{e \in E} \alpha(e) - \sum_{e \in E_s^\alpha} x_s^\alpha + \sum_{e \in \cup_{i=1}^{s-1} E_i^\alpha} \rho.$$

Since α is an edge-distribution we can replace $\sum_{e \in E(G)} \alpha(e)$ with 1. Doing so in the above equation in addition to replacing ρ by its value gives us

$$\sum_{e \in E(G)} \gamma(e) = 1 - \left(x_s^\alpha |E_s^\alpha| - \frac{x_s^\alpha |E_s^\alpha|}{\sum_{i=1}^{s-1} |E_i^\alpha|} \sum_{i=1}^{s-1} |E_i^\alpha| \right) = 1.$$

By definition γ has exactly $s-1$ strictly positive weights and hence, γ is an edge-distribution and the induction assumption applies to γ . Let H' be a minimum connected spanning subgraph of γ . By the induction assumption $\gamma(H') \leq \text{opt}$. We next show that $\alpha(H) \leq \gamma(H')$ and hence $\alpha(H) \leq \text{opt}$. According to the construction of γ we have $x_i^\gamma > x_j^\gamma$ if and only if $x_i^\alpha > x_j^\alpha$ for any $i, j \in \{1, \dots, s-1\}$ and therefore $E_i^\alpha = E_i^\gamma$ for $i = 1, \dots, s-1$. This in turn implies that $cr_i^\alpha = cr_i^\gamma$ for $i = 1, \dots, s-1$. According to Fact 2 we have

$$\gamma(H') = \sum_{i=1}^{s-1} x_i^\gamma cr_i^\gamma |E_i^\gamma|.$$

By replacing E_i^α with E_i^γ and cr_i^α with cr_i^γ and x_i^γ with $x_i^\alpha + \rho$ for $i = 1, \dots, s-1$ we get

$$\gamma(H') = \sum_{i=1}^{s-1} (x_i^\alpha + \rho) cr_i^\alpha |E_i^\alpha|.$$

This implies

$$\gamma(H') = \sum_{i=1}^s x_i^\alpha cr_i^\alpha |E_i^\alpha| - x_s^\alpha cr_s^\alpha |E_s^\alpha| + \rho \sum_{i=1}^{s-1} cr_i^\alpha |E_i^\alpha|.$$

By Fact 2, we can replace $\sum_{i=1}^s x_i^\alpha cr_i^\alpha |E_i^\alpha|$ by $\alpha(H)$. By also replacing ρ with its value we have

$$\gamma(H') = \alpha(H) - x_s^\alpha |E_s^\alpha| \left(cr_s^\alpha - \frac{\sum_{i=1}^{s-1} cr_i^\alpha |E_i^\alpha|}{\sum_{i=1}^{s-1} |E_i^\alpha|} \right). \quad (10)$$

Now by (8) and (10) we have $\alpha(H) \leq \gamma(H')$.

Assume that $\alpha(H) = \text{opt}$. Since $\alpha(H) \leq \gamma(H') \leq \text{opt}$ we get $\gamma(H') = \text{opt}$. Thus by the induction assumption γ is strong and hence $cr_i^\gamma = \text{opt}$ for $i = 1, \dots, s-1$. Since $cr_i^\alpha = cr_i^\gamma = \text{opt}$ for $i = 1, \dots, s-1$, to conclude that α is strong. Thus, we only need to show that $cr_s^\alpha = \text{opt}$. By replacing $\alpha(H), \gamma(H'), cr_1^\alpha, \dots, cr_{s-1}^\alpha$ with opt in (10) we get

$$cr_s^\alpha = \frac{\sum_{i=1}^{s-1} \text{opt} |E_i^\alpha|}{\sum_{i=1}^{s-1} |E_i^\alpha|} = \text{opt}.$$

(b) Let

$$\chi = (x_{s-1}^\alpha - x_s^\alpha) \left(1 + \frac{|E_s^\alpha|}{\sum_{i=1}^{s-1} |E_i^\alpha|}\right)^{-1}.$$

Let δ be such that

$$\delta(e) = \begin{cases} \alpha(e) + \chi & \text{if } e \in E_s^\alpha, \\ \alpha(e) - \chi \frac{|E_s^\alpha|}{\sum_{i=1}^{s-1} |E_i^\alpha|} & \text{if } e \in \cup_{i=1}^{s-1} E_i^\alpha, \\ 0 & \text{otherwise.} \end{cases} \quad (11)$$

Note that χ is such that $x_s^\alpha + \chi = x_{s-1}^\alpha - \chi \frac{|E_s^\alpha|}{\sum_{i=1}^{s-1} |E_i^\alpha|}$. Consequently, δ assigns the same weight to each edge in $E_{s-1}^\alpha \cup E_s^\alpha$ and hence δ has exactly $s-1$ strictly positive weights.

We next show that δ is an edge-distribution. By definition

$$\sum_{e \in E(G)} \delta(e) = \sum_{e \in E(G)} \alpha(e) + \sum_{e \in E_s^\alpha} \chi - \sum_{e \in \cup_{j=1}^{s-1} E_j^\alpha} \chi \frac{|E_s^\alpha|}{\sum_{i=1}^{s-1} |E_i^\alpha|}.$$

Since α is an edge-distribution we can replace $\sum_{e \in E(G)} \alpha(e)$ by 1 in the above to conclude

$$\sum_{e \in E(G)} \delta(e) = 1 + \chi \left(|E_s^\alpha| - \frac{|E_s^\alpha|}{\sum_{i=1}^{s-1} |E_i^\alpha|} \sum_{i=1}^{s-1} |E_i^\alpha| \right) = 1.$$

Note that by the choice of χ we have $\delta(e) > x_s^\alpha$ for every $e \in \cup_{i=1}^s E_i^\alpha$ and since $\delta(e) = 0$ for any other edge all the weights δ assigns are non-negative. Thus δ is an edge-distribution. Let H' be a minimum connected spanning subgraph of δ . Now as δ is an edge-distribution with $s-1$ strictly positive weights, by the induction assumption we have $\delta(H') \leq \text{opt}$. We conclude the claim by showing that $\alpha(H) < \delta(H')$.

By Fact 2 we have

$$\delta(H') = \sum_{i=1}^{s-1} x_i^\delta cr_i^\delta |E_i^\delta|. \quad (12)$$

According to the construction of δ we have $x_i^\delta > x_j^\delta$ if and only if $x_i^\alpha > x_j^\alpha$ for any $i, j \in \{1, \dots, s-2\}$ and therefore $E_i^\alpha = E_i^\delta$ for $i = 1, \dots, s-2$ which in turn implies that $cr_i^\alpha = cr_i^\delta$ for $i = 1, \dots, s-2$. Thus $E_i^\alpha = E_i^\delta$ and $cr_i^\alpha = cr_i^\delta$ for $i = 1, \dots, s-2$. Consequently by replacing $|E_{s-1}^\delta|$ with $|E_{s-1}^\alpha| + |E_s^\alpha|$ and cr_i^γ with cr_i^α for $i = 1, \dots, s-2$ in (12) we get

$$\delta(H') = \sum_{i=1}^{s-2} x_i^\delta cr_i^\alpha |E_i^\alpha| + x_{s-1}^\delta cr_{s-1}^\delta (|E_{s-1}^\alpha| + |E_s^\alpha|). \quad (13)$$

By definition of cut-rate we have

$$cr_{s-1}^\delta = \frac{|E_{s-1}^\alpha| cr_{s-1}^\alpha + |E_{s-1}^\alpha| cr_s^\alpha}{|E_{s-1}^\alpha| + |E_s^\alpha|},$$

Plugging this in (13) gives us

$$\delta(H') = \sum_{i=1}^{s-2} x_i^\delta cr_i^\alpha |E_i^\alpha| + x_{s-1}^\delta cr_{s-1}^\alpha |E_{s-1}^\alpha| + x_{s-1}^\delta cr_s^\alpha |E_s^\alpha|.$$

Since $x_s^\delta = x_s^\alpha + \chi$ and $x_i^\delta = x_i^\alpha - \chi \frac{|E_s^\alpha|}{\sum_{i=1}^{s-1} |E_i^\alpha|}$ for $i = 1, \dots, s-2$ we get

$$\delta(H') = \sum_{i=1}^s x_i^\alpha cr_i^\alpha |E_i^\alpha| + \chi |E_s^\alpha| \left(cr_s^\alpha - \frac{\sum_{i=1}^{s-1} cr_i^\alpha |E_i^\alpha|}{\sum_{i=1}^{s-1} |E_i^\alpha|} \right). \quad (14)$$

By Fact 2 we can also replace $\sum_{i=1}^s x_i^\alpha cr_i^\alpha |E_i^\alpha|$ with $\alpha(H)$ in (14). This together with (9) implies that $\delta(H') > \alpha(H)$. Note that as $opt > \delta(H')$ it can not be the case that $\alpha(H) = opt$. \square

2 Appendix: Prime partition

In this section we prove Theorem 2. To do so, we introduce a polynomial time algorithm that on input graph $G = (V, E)$ returns a partition of E , which afterwards we show is the prime-partition of G . The algorithm uses oracle access to a routine **PrimeSet** that given a graph returns its cut-rate and one of its prime sets. This routine runs in time polynomial in the size of G and is introduced in Subsection 2.3

2.1 Construction of the prime-partition

Algorithm 1 Prime-partition construction

Input: Graph G .

Output: Prime partition \mathcal{P} .

```

1:  $\mathcal{P} \leftarrow \emptyset$ 
2: if  $E(G) = \emptyset$  then
3:   return  $\mathcal{P}$ 
4: end if
5:  $i \leftarrow 1$ 
6:  $(opt, P_i) \leftarrow \text{PrimeSet}(G)$ 
7:  $G_i \leftarrow G$ 
8: repeat
9:    $\mathcal{P} \leftarrow \mathcal{P} \cup \{P_i\}$ 
10:   $G_{i+1} \leftarrow G_i \setminus P_i$ 
11:   $i \leftarrow i + 1$ 
12:   $(c, P_i) \leftarrow \text{PrimeSet}(G_i)$ 
13: until  $c < opt$ 
14: if  $E(G_i) \neq \emptyset$  then
15:    $\mathcal{P} \leftarrow \mathcal{P} \cup \{P_i\}$ 
16: end if
17: return  $\mathcal{P}$ 

```

Each computation done by Algorithm 1 requires running time polynomial in $|V|$ including the calls to **PrimeSet** according to Appendix 2.3. Therefore, the only reason the running time of Algorithm 1 may be too long is the repeat loop. Note that, if **PrimeSet** returns an empty set, then it also sets $c = 0$. Thus, after any iteration of the repeat loop that is not the last the number of edges of G' is decreased by at least 1. Observe that **PrimeSet** returns $(0, \emptyset)$ and does not reach the repeat loop if $opt = 0$.

Thus only if $\text{opt} > 0$ then the repeat loop is reached and then the above ensures that it goes through at most $|E|$ iterations. Hence the running time of Algorithm 1 is polynomial in the size of G .

From here on in this section $t = |\mathcal{P}|$, where \mathcal{P} is the output of Algorithm 1 on input graph G , and the elements of \mathcal{P} are named as they were named by Algorithm 1, thus $\mathcal{P} = \{P_1, \dots, P_t\}$. In addition let $E^0 = \emptyset$ and for each $k = 1, \dots, t$ let $E^k = \cup_{i=1}^k P_i$ and r_k be the cut-rate of P_k in $G \setminus E^{k-1}$.

2.2 The output of Algorithm 1 is the prime partition

Proposition 8. *There exists a maxmin-edge-distribution β such that $\mathcal{E}(\beta) = \mathcal{P}$.*

Proof. Set $\rho = \frac{1}{\sum_{i=1}^t (t-i)|E_i|}$ and let $\beta : E \rightarrow \mathbb{R}$, where for each $i = 1, \dots, t$ and $e \in P_i$ we have $\beta(e) = (t-i)\rho$. Observe that

$$\sum_{e \in E} \beta(e) = \sum_{i=1}^t x_i^\beta |E_i^\beta| = \sum_{i=1}^t \rho(t-i) |E_i^\beta| = \rho \sum_{i=1}^t (t-i) |E_i^\beta| = \rho \rho^{-1} = 1$$

and hence β is an edge-distribution. By definition $E_i^\beta = P_i$ for $i = 1, \dots, t$. We next show that $r_i = \text{opt}$ for $i = 1, \dots, t-1$. By Theorem 1 this implies that β is a maxmin-edge-distribution.

Algorithm 1 selects P_1 so that $r_1 = \text{opt}$. Let $k < t$ and assume that $r_j = \text{opt}$ for every $j < k$. Hence by Fact 3 we have $\text{cr}(E^{k-1}) = \text{opt}$. Consequently $r_k \leq \text{opt}$ since otherwise by Fact 3 we get that $\text{cr}(E^{k-1}) > \text{opt}$. As P_k is not the last set added to \mathcal{P} by Algorithm 1, we have $r_k \geq \text{opt}$ and hence it is the case that $r_k = \text{opt}$. \square

We next show that \mathcal{P} refines $\mathcal{E}(\alpha)$ for every maxmin-edge-distribution α . We start with a simple case that we use later on to prove the general result.

Proposition 9. *If $\text{cr}(E') = \text{opt}$ for $E' \subseteq E$ then \mathcal{P} refines $\{E', E \setminus E'\}$.*

Proof. If $E' = E$ then the proposition holds trivially. Hence we only need to prove the proposition holds when $E' \subset E$. We show that $P_i \cap E' = \emptyset$ or $P_i \subseteq E'$ for every $i = 1, \dots, t$. Let $k \in \{1, \dots, t\}$. If $E' \setminus E^{k-1} = \emptyset$ then $P_k \cap E' = \emptyset$. Therefore we only need to deal with the case that $E' \setminus E^{k-1} \neq \emptyset$. Assume this is indeed so. By Proposition 8, we have $r_i = \text{opt}$ for $i = 1, \dots, t-1$ hence by Fact 3 we get $\text{cr}(E^{k-1}) = \text{opt}$. Since also $\text{cr}(E') = \text{opt}$, by Proposition 7, we have $\text{opt}_{G \setminus E^{k-1}} = \text{opt}$ and $\text{cr}_{G \setminus E^{k-1}}(E' \setminus E^{k-1}) = \text{opt}$. We separate the proof into two cases the first $k = 1, \dots, t-1$ and in the second $k = t$.

Recall that Algorithm 1 selects P_k so that it is a prime-set in $G \setminus E^{k-1}$. So now $\text{opt}_{G \setminus E^{k-1}} = \text{opt}$ and $\text{cr}_{G \setminus E^{k-1}}(P_k) = \text{cr}_{G \setminus E^{k-1}}(E' \setminus E^{k-1}) = \text{opt}$. Thus by Proposition 7 either $P_k \cap (E' \setminus E^{k-1}) = \emptyset$ or $P_k \subseteq (E' \setminus E^{k-1})$. If $P_k \subseteq (E' \setminus E^{k-1})$ then $P_k \subseteq E'$, and if $P_k \cap (E' \setminus E^{k-1}) = \emptyset$ then $P_k \cap E^{k-1} = \emptyset$ because $P_k \cap E^{k-1} = \emptyset$.

Assume $k = t$ and for the sake of contradiction that $E' \setminus E^{t-1} \neq \emptyset$. Since we have shown that $\text{opt}_{G \setminus E^{t-1}} = \text{opt}$ and $\text{cr}_{G \setminus E^{t-1}}(E' \setminus E^{t-1}) = \text{opt}$ it is the case that $G \setminus E^{t-1}$ has a prime-set that has cut-rate opt in $G \setminus E^{t-1}$. This subset is strictly contained in $E \setminus E^{t-1}$ Proposition 7 implies that every prime set in $E \setminus E^{t-1}$ is strictly contained in $E \setminus E^{t-1}$. Hence, Algorithm 1 would have found a prime-set $E^* \subset E \setminus E^{t-1}$ and added it to \mathcal{P} . That is $E^* \in \mathcal{P}$. Yet this can not be since by construction \mathcal{P} is a partition of E . \square

Proposition 10. *If γ is a maxmin-edge-distribution then \mathcal{P} refines $\mathcal{E}(\gamma)$.*

Proof. Let $t = |\mathcal{E}(\gamma)|$. Recall that since γ is a maxmin-edge-distribution by definition for $i = 1, \dots, t-1$ we have $cr_i^\gamma = opt$. If $t = 1$ then the only set in $\mathcal{E}(\gamma)$ is E and hence the lemma trivially holds. By Proposition 9 the lemma also holds when $|\mathcal{E}(\gamma)| = 2$. Assume by way of induction that proposition holds for any partition $\mathcal{S} = \{E_1, \dots, E_{t-1}\}$ of E such that $cr_{G \setminus \cup_{i=1}^{\ell} E_i}(E_\ell) = opt$. Let $\mathcal{S}_1 = \{\cup_{i=1}^{t-1} E_i^\alpha, E_t^\alpha\}$ and $\mathcal{S}_2 = \{E_1, \dots, E_{t-2}, E_{t-1}^\alpha \cup E_t^\alpha\}$. Note that if \mathcal{P} refines both \mathcal{S}_1 and \mathcal{S}_2 then it refines \mathcal{S} . By the induction assumption, \mathcal{P} refines \mathcal{S}_2 . By Fact 3 we have that $cr(\cup_{i=1}^{m-1} E_i) = opt$ and therefore \mathcal{P} refines \mathcal{S}_1 by Proposition 9. \square

2.3 PrimeSet

In this section we explain the subroutine for finding a minimal optimal set, which we call a **PrimeSet**. We assume that the graph is connected, in case it is not connected we run the routine separately on each connected component and return the **PrimeSet** (and value opt that achieves the largest opt among these connected components. If there is more than one, pick one arbitrarily. By Fact 3, the cut-rate of $\cup_{i=1}^{m-1} E_i$ in G is opt and therefore by Proposition 9, we have that \mathcal{P} refines \mathcal{S}_1 . For our goal, we extend the notion of cut-rate of a graph to edge weighted graphs.

Definition 16. Let $E' \subseteq E$ and $\omega : E \rightarrow \mathbb{R}^+$. The cut-rate of E' in G, ω is denoted by $cr_\omega(E')$ and defined as follows.

$$cr_\omega(E') := \begin{cases} \frac{C_G(E') - C_G}{\omega(E')} & \text{if } |V| > 1 \text{ and } |E'| > 0, \\ 0 & \text{otherwise.} \end{cases} \quad (15)$$

The cut-rate of G, ω where $\omega : E \rightarrow \mathbb{R}^+$ is defined as

$$opt_\omega := \max_{E' \subseteq E} cr_\omega(E') \quad (16)$$

There exists strongly polynomial algorithms in [5, 26, 4] that on G, ω returns opt_ω . We shall assume from here on that opt_ω is given and omit the fact that this is done by the mentioned algorithm.

A prime-set of G is found as follows. If $E = \emptyset$ then stop and return $(0, \emptyset)$. Otherwise, set $\omega : E \rightarrow \mathbb{R}^+$ so that $\omega(e) = 1$ for every $e \in E$. Note that in this case $opt = opt_\omega$ and hence we assume opt is known. Set $\omega' = \omega$. Next iterate e over the elements of E according to some arbitrary order and in each iteration do the following. Set $\omega'(e)$ to be 2 and if $opt_{\omega'} = opt$ then set ω to be ω' and otherwise set ω' to be ω . That is, $\omega(e)$ is changed only if $opt_{\omega'} = opt$ and otherwise remains the same. After the iterative process is over set $E' = \{e \in E \mid \omega(e) = 1\}$ and return (opt, E') . Note that the total number of operations done is polynomial in the size of G and so is the running time. To show that indeed this achieves our goal, we only need to prove that E' is a prime set of G .

Let us look at any fixed iteration over e . By definition ω is changed only if $opt_{\omega'} = opt$ and then it is set to ω' . This implies that there exists $E^* \subseteq E$ such that $cr_{\omega'}(E^*) = opt$. Now it can not be the case that $\omega'(e) = 2$ for some $e \in E^*$, since this would imply that $cr(E^*) > opt$. Consequently, $cr_\omega(E^*) = opt$. This is true for any fixed iteration and hence also for the last. Therefore, there exists $E'' \subseteq E'$ such that $cr(E'') = opt$. Finally assume for the sake of contradiction that $E'' \subset E'$. Let $e' \in E' \setminus E''$. This implies that in the iteration dedicated to e' we had $opt_{\omega'} \neq opt$. Since at this stage $\omega'(e) = 1$ for every E'' it must be that $opt_{\omega'} > opt$. Yet this can not be since the weights assigned to each edge by ω' is at least as that assigned by ω and hence at every iteration $opt_{\omega'} \leq opt$.

3 Appendix: Proof of Lemma 1

If $cr(E) = opt$, then an edge-distribution that assigns equal weight to all edges is a maxmin-edge-distribution and so there is no degenerate set.

We now prove that if $cr(E) \neq opt$, then the degenerate set exists, it is unique, and can be found in running time polynomial in the size of G .

Assume that $cr(E) \neq opt$. By the definition of opt , this can only happen if $cr(E) < opt$. Let β be a maxmin-edge-distribution such that $\mathcal{E}(\beta) = \mathcal{P}$ and set $t = |\mathcal{P}|$. By definition, E_t^β assigns strictly positive weights to the edges in each E_i^β for every $i = 1, \dots, t-1$. Hence, the only candidate for being the degenerate set is E_t^β . We next show that this is indeed the case.

Assume for the sake of contradiction that there exists a maxmin-edge-distribution γ that assigns strictly positive weights to the edges in E_t^β . Let $d = \min_{i \in \{1, \dots, t-1\}} \{x_i^\beta - x_{i-1}^\beta\}/10$ and set $\delta = (1-d)\beta + d\gamma$ (the choice of 10 is arbitrary). Observe that δ has the same number of distinct weights as γ , and $\mathcal{E}(\delta) = \mathcal{E}(\beta)$ and we have $E_t^\delta = E_t^\beta$ for $i = 1, \dots, t$. Let H be a minimum connected spanning subgraph of δ . Since δ is a convex combination of maxmin-edge-distributions it is a maxmin-edge-distribution and therefore, by Corollary 1, we have $\delta(H) = opt$. We next get the required contradiction by showing that $\delta(H) < opt$.

Since $\mathcal{E}(\delta)$ is a partition of E we have

$$\delta(H) = \sum_{i=1}^t x_i^\delta |E(H) \cap E_i^\delta|. \quad (17)$$

As δ is a maxmin-edge-distribution, by Theorem 1, we have $cr_i^\delta = opt$ for $i = 1, \dots, t-1$ and hence $cr_t^\delta < opt$, since otherwise, by Fact 3, we have $cr(E) \geq opt$. By Proposition 1, we have $|E(H) \cap E_i^\delta| = opt|E_i^\delta|$ for $i = 1, \dots, t-1$. Applying this to (17) we get

$$\delta(H) = x_t^\delta |E(H) \cap E_t^\delta| + opt \sum_{i=1}^{t-1} x_i^\delta |E_i^\delta|. \quad (18)$$

Now, since H is a minimum connected spanning subgraph, $|E(H) \cap E_t^\delta|$ is the minimum possible, which in this case is $cr_t^\delta |E_t^\delta|$. Since $cr_t^\delta < opt$, we get that $cr_t^\delta |E_t^\delta| < opt|E_t^\delta|$. Thus, by replacing $|E(H) \cap E_t^\delta|$ by $opt|E_t^\delta|$ in (18), we get

$$\delta(H) < opt \sum_{i=1}^t x_i^\delta |E_i^\delta| = opt,$$

where the equality is because δ is an edge-distribution.

We now explain how to compute D . Once opt is known, one only needs to check if $\frac{|V|-1}{|E|} = opt$. If the answer is yes, then there is no degenerate set; if the answer is no, then D is the last set inserted to \mathcal{P} by Algorithm 1.

4 Appendix: Proof of Proposition 3

Let H be an omni-connected-spanning-subgraph and α such that \mathcal{P} refines $\mathcal{E}(\alpha)$ and $\alpha(e) = 0$ for every $e \in D$. Thus for each $P \in \mathcal{P}$ there exists y^P such that $\alpha(e) = y^P$ for every $e \in P$. Since \mathcal{P} is a partition of E , we have $\alpha(H) = \sum_{P \in \mathcal{P}} y^P |H \cap P|$ and

as H is an omni-connected-spanning-subgraph also $|H \cap P| = |P|opt$ for every $P \in \mathcal{P}$. Consequently,

$$\alpha(H) = \sum_{P \in \mathcal{P}} y^P |P|opt = opt \sum_{P \in \mathcal{P}} y^P |P|.$$

Now, as α is an edge-distribution and $\sum_{P \in \mathcal{P}} y^P |P| = 1$, with the above, we get $\alpha(H) = opt$. Now, if α is maxmin-edge-distribution, by Corollary 1, the value of the game is opt and H is a minimum connected spanning subgraph of α . \square

5 Appendix: Partial Order

5.1 Proof of Theorem 3

Note that if $\mathcal{P} = 1$ then the theorem trivially holds hence we assume $\mathcal{P} > 1$.

Let α be a maxmin-edge-distribution. By Lemma 1 we have $\alpha(e) = 0$ for every $e \in D$. Assume for the sake of contradiction that α does not agree with \mathcal{R} . By Definition 7, we have that \mathcal{P} refines α and hence since α does not agree with \mathcal{R} there exist $P \in \mathcal{P} \setminus \{D\}$ that leads to $P' \in \mathcal{P} \setminus \{D\}$, an omni-connected-spanning-subgraph H , $e \in E(H) \setminus P$ and $e' \in P' \cap E(H)$ such that $\alpha(e) < \alpha(e')$ and $H' = (H \setminus \{e\}) \cup \{e'\}$ is a connected spanning subgraph. Since H is an omni-connected-spanning-subgraph and α a maxmin-edge-distribution by Proposition 3 we have $\alpha(H) = opt$. Thus $\alpha(H') = \alpha(H) + (\alpha(e) - \alpha(e')) < opt$. This is in contradiction to Fact 1, which implies that $\alpha(H') \geq opt$ since α is a maxmin-edge-distribution.

Assume α is an edge-distribution that agrees with \mathcal{R} . We next show that this implies that α is a maxmin-edge-distribution.

Let m be the number of the strictly positive weights of α . Assume by way of contradiction that α is not a maxmin-edge-distribution. By Theorem 1 this can only happen if there exists $i \in \{1, \dots, m\}$ such that $cr_i^\alpha \neq opt$. Let ℓ be the smallest element in $\{1, \dots, m\}$ such that $cr_\ell^\alpha \neq opt$. Let $E' = \cup_{i=1}^\ell E_i^\alpha$. We show next that $cr(E') < opt$. If $\ell = 1$ then $cr_\ell^\alpha \leq opt$ since cr_ℓ^α is the cut-rate of E_ℓ in G . Thus in this case $E' = E_\ell$, and the goal is achieved. Assume that $\ell > 1$. By the minimality of ℓ , we have that $cr_i^\alpha = opt$ for $i = 1, \dots, \ell - 1$. If $cr_\ell^\alpha > opt$ then by Fact 3 we have $cr(\cup_{i=1}^\ell E_i^\alpha) > opt$ which is a contradiction to the definition of opt . Thus, as $cr_\ell^\alpha \neq opt$ we have $cr_\ell^\alpha < opt$ and hence again by Fact 3 the $cr(E') < opt$.

Let C_1, \dots, C_s be the connected components of $G \setminus E'$. Let H be an omni-connected-spanning-subgraph and let H_1, \dots, H_r be the connected components of $E(H) \setminus E'$. Note that for each $i \in \{1, \dots, r\}$ there exists a unique $j \in \{1, \dots, s\}$ such that $E(H_i) \subseteq E(C_j)$. For each $j \in \{1, \dots, s\}$ set I_j to be the set of all $i \in \{1, \dots, r\}$ such that $E(H_i) \subseteq E(C_j)$. Assume that $s < r$, we shall show afterwards that this is indeed true. By the pigeon-hole principle there exists $j \in \{1, \dots, r\}$ such that $|I_j| > 1$. Since C_j is a connected component and H a connected spanning subgraph of G there exist $x, y \in I_j$ and $e = \{u, v\} \in E(C_j) \setminus \cup_{i=1}^{|I_j|} E(H_i)$ such that $u \in V(H_x)$ and $v \in V(H_y)$. Since H is a connected spanning subgraph there is a path in H between u and v this path contains edges not in $E(C_j)$ since u, v are in different connected components of $H \setminus E'$. Thus this path contains an edge $e' \in E'$ since only edges from E' connect the vertices of C_j to the rest of the graph. Consequently $(H \setminus \{e\}) \cup \{e'\}$ is a connected spanning subgraph of G . Let $P, P' \in \mathcal{P}$ be such that $e \in P$ and $e' \in P'$. By the above P leads to P' . Yet, this can not be since $\alpha(e) < \alpha(e')$ and α agrees with \mathcal{R} .

It remains to show that indeed $s < r$. By the definition of cut-rate the number of connected components s in $G \setminus E'$ is $cr(E')|E'|$, which is strictly less than $opt|E'|$. Now as α agrees with \mathcal{R} we know that \mathcal{P} refines E' . Hence E' is the union of sets

in $\mathcal{P} \setminus \{D\}$. Consequently, by the definition of a omni-connected-spanning-subgraph, we have $E(H) \cap E' = \text{opt}|E'|$. Hence $r = \text{opt}|E'|$ because the number of connected components in $H \setminus E'$ is the number of edges in $E(H) \cap E'$. \square

5.2 Proof of Theorem 4

Definition 17. We say that $P \in \mathcal{P}$ is an ancestor of $P' \in \mathcal{P}$ if there is a chain in \mathcal{O} from P to P' .

We show that for each $P \in \mathcal{P} \setminus \{D\}$ we can find all of the ancestors of P . Once we know the ancestors of each element $\mathcal{P} \setminus \{D\}$ finding the parent of each such element is easy. An ancestor P of P' is also a parent of P' if there does not exist an P^* , that is neither P nor P' , such that P is an ancestor of P^* and P^* is an ancestor of P' . Checking this for each pair element and each one of its ancestors requires running time that is polynomial in the size of G .

To achieve our goal we need the following proposition.

Proposition 11. Let $\mathcal{P}^* \subseteq \mathcal{P} \setminus \{D\}$, and $P^* \in \mathcal{P}^*$ and set $E^* = \cup_{P \in \mathcal{P}^*} P$ then

- $cr(E^*) = \text{opt}$ if \mathcal{P}^* contains only P^* and all its ancestors.
- If $cr(E^*) = \text{opt}$ and \mathcal{P}^* contains an element that is not P^* or one of its ancestors then it also contains such a P for which $cr(E^* \setminus P) = \text{opt}$.

Proof. Set $\alpha : E(G) \rightarrow \mathbb{R}$ so that $\alpha(e) = \frac{1}{|E^*|}$ if $e \in E^*$ and $\alpha(e) = 0$ otherwise. Now α is an edge-distribution, \mathcal{P} refines $\mathcal{E}(\alpha)$ and $E_1^\alpha = E^*$, $E_2^\alpha = E \setminus E^*$ and $\alpha(e) = 0$ for every $e \in D$.

We now prove the first item. For any parent and its child if the child is in \mathcal{P}^* it is either P^* or one of its ancestors. Thus the parent is also an ancestor of P^* and hence is also in \mathcal{P}^* . Consequently α agrees with \mathcal{O} and hence by Theorem 3, we have that α is a maxmin-edge-distribution. This in turn by Theorem 1 implies $cr(E^*) = \text{opt}$.

We now prove the second item. Assume $cr(E^*) = \text{opt}$ and \mathcal{P}^* contains an element that is not P^* or one of its ancestors. Then there exists $P \in \mathcal{P}^*$ that is not an ancestor of any other element in \mathcal{P}^* . Since $cr(E^*) = \text{opt}$ by Theorem 1 α is a maxmin-edge-distribution. Hence by Theorem 1 α agrees with \mathcal{O} .

Set $\beta : E(G) \rightarrow \mathbb{R}$ so that $\beta(e) = \frac{1}{|E^*|}$ if $e \in E^*$ and $\beta(e) = 0$ otherwise. Now β is an edge-distribution, \mathcal{P} refines $\mathcal{E}(\beta)$ and $E_1^\beta = E^*$, $E_2^\beta = E \setminus E^*$ and $\beta(e) = 0$ for every $e \in D$. The only way that β does not agree with \mathcal{O} is if a child of P is in $\mathcal{P}^* \setminus \{P\}$, yet this can not be, since P is not an ancestor of any element in \mathcal{P}^* . Thus, β agrees with \mathcal{O} and hence, by Theorem 1, β is a maxmin-edge-distribution. By Theorem 1, this implies that $cr(E^* \setminus P) = \text{opt}$. \square

We next show how to find the ancestors of P' . Set $\mathcal{P}' = \mathcal{P} \setminus \{D\}$ and $E' = \cup_{P \in \mathcal{P}'} P$. If there exists $P^* \in \mathcal{P}'$ such that $P^* \neq P'$ and $cr(E' \setminus P) = \text{opt}$ remove it from \mathcal{P}' and recompute E' . Repeat until no such element is found.

Note that this requires $|V|$ repetitions each taking a polynomial time in the size of G . Consequently, the running time is polynomial in the size of G .

When $\mathcal{P} = \mathcal{P} \setminus \{D\}$ we have $cr(E') = \text{opt}$ because of the following. By definition there exists a maxmin-edge-distribution β such that $\mathcal{E}(\beta) = \mathcal{P}$. Note that \mathcal{P}' is all the non-degenerate sets in $\mathcal{E}(\beta)$ and hence $cr(E') = cr(\cup_{i=1}^m E_i^\beta)$. By Theorem 1 $cr_i^\beta = \text{opt}$ for $i = 1, \dots, m$, where m is the maximal index such that $x_m^\beta > 0$. Hence according to Fact 3 we have $cr(\cup_{i=1}^m E_i^\beta) = \text{opt}$.

Finally we show that at the end what remains in \mathcal{P}' is only P' and all its ancestors. The set P' is never removed from \mathcal{P}' . By Proposition 11 for any ancestor P^* of P'

it is the case that $cr(E' \setminus P^*) < opt$ and hence none of the ancestors of P' are ever removed. Also by Proposition 11 as long as \mathcal{P}' does not contain only P' and each one of its ancestors there exists a P^* such that $cr(E' \setminus P^*) = opt$ and hence such an element will be removed. Thus only P' and each one of its ancestors are never removed and consequently they are the only elements remaining in \mathcal{P}' at the end of the process. \square

6 Appendix: Proof of Lemma 2

Let β be a maxmin-edge-distribution such that one of the following holds

1. There exists $P \in \mathcal{P} \setminus \{D\}$ that is a parent of $P' \in \mathcal{P} \setminus \{D\}$ such that $\beta(e) = \beta(e')$ for every $e \in P$ and $e' \in P'$.
2. There exist $P \in \mathcal{P} \setminus \{D\}$ such that $\beta(e) = 0$ for every $e \in P$.

We shall show that β has a minimum connected spanning subgraph that is not an omni-connected-spanning-subgraph. Afterwards we shall show that for every γ for which both conditions do not hold, every minimum connected spanning subgraph of γ is an omni-connected-spanning-subgraph. According to Proposition 3, every omni-connected-spanning-subgraph is a minimum connected spanning subgraph of γ , this means that such γ are the only maxmin-edge-distributions that have the minimum possible number of minimum connected spanning subgraphs.

Assume the first condition holds for β . By the Definition 11 there exists an omni-connected-spanning-subgraph H and edges $e_1 \in P \setminus E(H)$, $e_2 \in P' \cap H$ such that $H' = (T \cup \{e_1\}) \setminus \{e_2\}$ is a connected spanning subgraph. Observe that H' is not an omni-connected-spanning-subgraph of β but is a minimum connected spanning subgraph of β since $\beta(H') = \beta(H) = opt$. Assume the second condition holds for β . Let H be an omni-connected-spanning-subgraph. Recall we assumed $opt < 1$ and hence as H is an omni-connected-spanning-subgraph we have $|E(H) \cap P| = opt|P| < |P|$ and therefore there exists $e \in P \setminus E(H)$. Since H is a minimum connected spanning subgraph of β by Proposition 3 and $\beta(e) = 0$ we also have $H \cup \{e\}$ is a minimum connected spanning subgraph of β . Note that $H \cup \{e\}$ is not an omni-connected-spanning-subgraph.

Let γ be some maxmin-edge-distribution for which the above two conditions do not hold. That is, $\gamma(e) > 0$ for every $e \in E \setminus D$ and $P \in \mathcal{P} \setminus \{D\}$ that is a parent of $P' \in \mathcal{P} \setminus \{D\}$ and every $e \in P$, $e' \in P'$ we have $\gamma(e) > \gamma(e')$.

From here on let H be a minimum connected spanning subgraph of γ . We next show that H is an omni-connected-spanning-subgraph. Let m be the number of distinct strictly positive values of γ and set $\mathcal{P}_i = \{P \in \mathcal{P} \mid P \in E_i^\gamma\}$ for $i = 1, \dots, m$. Note that by the definition of γ for every $P \in \mathcal{P} \setminus \{D\}$ there exists $i \in \{1, \dots, m\}$ such that $P \in \mathcal{P}_i$. Assume by way of contradiction that H is not an omni-connected-spanning-subgraph. Let k be the minimum integer such that there exists $P \in \mathcal{P}_k$ for which $|H \cap P| \neq |P|opt$. Since H is a minimum connected spanning subgraph and γ a maxmin-edge-distribution by Proposition 1 we have $|H \cap E_k^\gamma| = opt|E_k^\gamma|$. Since \mathcal{P} refines $\mathcal{E}(\gamma)$ we also have $|H \cap E_k^\gamma| = \sum_{E' \in \mathcal{P}_k} |H \cap E'|$ and $|E_k^\gamma| = \sum_{E' \in \mathcal{P}_k} |E'|$ and hence

$$\sum_{E' \in \mathcal{P}_k} |H \cap E'| = opt \sum_{E' \in \mathcal{P}_k} |E'|$$

Therefore the fact that $|H \cap P| \neq |P|opt$ implies that there exists $P' \in \mathcal{P}_k$ such that $|H \cap P'| < |P'|opt$. Let P' be such a set.

Let E^* be the union of P' and all its ancestors (see Definition 17 in Section 5). We next show that $cr(E^*) = opt$. Set $\alpha : E(G) \rightarrow \mathbb{R}$ so that $\alpha(e) = \frac{1}{|E^*|}$ if $e \in E^*$ and $\alpha(e) = 0$ otherwise. Observe that α is an edge-distribution, \mathcal{P} refines $\mathcal{E}(\alpha)$ and

$E_1^\alpha = E^*$, $E_2^\alpha = E \setminus E^*$ and $\alpha(e) = 0$ for every $e \in D$. Now for any parent and its child if the child is contained E^* it is either P' or one of its ancestors. Thus the parent is also an ancestor of P' and hence is also in E^* . Consequently, α agrees with \mathcal{O} and hence by Corollary 2 we have that α is a maxmin-edge-distribution. This in turn by Theorem 1 implies $cr(E^*) = opt$.

Note that because of the strict weight inequalities, all the ancestors of P' are elements in one of the sets $\mathcal{P}_1, \dots, \mathcal{P}_{k-1}$. Thus for any ancestor P^* of P' we have $|H \cap P^*| = opt|P^*|$. Consequently, $|H \cap E^*| < opt|E^*|$ yet this can not be since $opt|E^*|$ is the minimum number of edges a connected spanning subgraph can have in E^* . \square

7 Appendix: The nucleolus

Proof. Let κ, ν be as stated in the theorem and H an omni-connected-spanning-subgraph and t the number of layers. Observe that

$$\sum_{e \in E} \nu(e) = \sum_{i=1}^t i \cdot |L_i| \cdot \kappa = \kappa \sum_{i=1}^t i \cdot |L_i| = \kappa \cdot \kappa^{-1} = 1$$

and hence ν is an edge-distribution. Note that by definition \mathcal{P} refines $\mathcal{E}(\nu)$ and for any $e \in P \in \mathcal{P}$ that is a parent of $e' \in P' \in \mathcal{P}$ we have $\nu(e) > \nu(e')$ and hence ν is a maxmin-edge-distribution and specifically a prime-edge-distribution.

We now show that the weight of the second best strategy of the hider is $opt + k$. Afterwards we show that the only ν is the only prime-edge-distribution for which the weight of the second best strategy of the hider is at least $opt + k$.

Let $P \in \mathcal{P}$ be such that $P \subseteq L_1$. Since $opt < 1$ Proposition 1 implies that there exists $e \in P \setminus E(H)$. By the definition of ν we have $\nu(e) = \kappa$. By Proposition 3 $\nu(H) = opt$ and hence $\nu(H \cup \{e\}) = opt + \kappa$. Note that since $\nu(e')$ is a multiple of κ for every $e' \in E$ there does not exist a connected spanning subgraph H' such that $opt < \nu(H') < opt + \kappa$.

Let α be a prime-edge-distribution such that the second smallest weight of a connected spanning subgraph is $opt + \kappa$. We shall prove by induction on ℓ that $\alpha(e) > \ell\kappa$ for every $e \in L_\ell$ and every $\ell = 1, \dots, t$. Since the only prime-edge-distribution that satisfies this conditions is ν this implies that $\alpha = \nu$.

Assume for the sake of contradiction that there exists $e \in P \in L_1$ such that $\alpha(e) < \kappa$. Since $opt < 1$ Proposition 1 implies that there exists $e' \in P \setminus E(H)$. By Proposition 3 we have $\nu(H) = opt$ and hence because $\alpha(e') = \alpha(e) < \kappa$ we get $\nu(H \cup \{e\}) = opt + \alpha(e') < opt + \kappa$. In addition as α is a prime-edge-distribution $\alpha(e') > 0$ and thus $\nu(H \cup \{e\}) > opt$. Yet the assumption was that the hiders second best response to α is at least $opt + \kappa$.

Assume by way of induction that for $\ell - 1$ we have $\alpha(e) > (\ell - 1)\kappa$ for every $e \in L_{\ell-1}$. Assume for the sake of contradiction that there exists $e \in P \in L_\ell$ such that $\alpha(e) < \ell \cdot \kappa$. By the definition of L_ℓ , there exists $P' \subseteq L_{\ell-1}$ such that P is a parent of P' . Consequently, there exists an omni-connected-spanning-subgraph H' , $e' \in P \setminus E(H')$ and $e'' \in P' \cap E(H')$ such that $H^* = (H' \setminus \{e''\}) \cup \{e'\}$ is a spanning tree of G . Observe that $\alpha(H^*) = \alpha(H') + \alpha(e') - \alpha(e'')$. By Proposition 3, we have $\alpha(H') = opt$ and hence $\alpha(H^*) = opt + \alpha(e') - \alpha(e'')$. By the induction assumption, we have $\alpha(e'') \geq (\ell - 1)\kappa$ and therefore as $\alpha(e') = \alpha(e) < \ell \cdot \kappa$ we get $\alpha(H^*) < opt + \kappa$. In addition as α is a prime-edge-distribution we have $\alpha(e') > \alpha(e'')$ and therefore $\nu(H^* \cup \{e\}) > opt$. Yet the assumption was that the hider's second best response to α is at least $opt + \kappa$. \square

8 Appendix: Extreme points

Definition 18. We say $\mathcal{B} \subseteq \mathcal{P} \setminus \{D\}$ is closed if for every $P \in \mathcal{B}$ all of the ancestors of P are also in \mathcal{B} . Set $A(\mathcal{B}) = \cup_{E' \in \mathcal{B}} E'$. We say a closed set \mathcal{B} is minimal if there do not exist any closed sets $\mathcal{B}_1, \mathcal{B}_2 \subset \mathcal{P} \setminus \{D\}$ such that $A(\mathcal{B}) = A(\mathcal{B}_1) \cup A(\mathcal{B}_2)$.

Definition 19. An edge-distribution α is an extreme-edge-distribution if and only if there exists a minimal closed set \mathcal{B} such that $E_1^\alpha = A(\mathcal{B})$ and $E_2^\alpha = 0$.

Theorem 6. The extreme-edge-distributions are the extreme points of the maxmin-polytope.

The proof uses similar techniques to our other proofs and is omitted.